

I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being transmitted via the Office electronic filing system in accordance with § 1.6(a)(4).

Dated: 12/5/2007
Electronic Signature for Warren S. Wolfeld: /Warren S. Wolfeld/

Docket No.: OPTI 3140-6US
(PATENT)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:
Subir K. Ghosh et al.

Application No.: 10/619,798

Confirmation No.: 7888

Filed: July 15, 2003

Art Unit: 2188

For: PREDICTIVE SNOOPING OF CACHE
MEMORY FOR MASTER-INITIATED
ACCESSES

Examiner: K. L. Ellis

MS Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**SECOND DECLARATION OF SUBIR GHOSH
UNDER 37 C.F.R. §1.131(b)**

I, SUBIR GHOSH, declare as follows:

1. I am one of the inventors named in the above-identified patent application. With this Declaration, I provide documentary evidence that the invention of the above-identified patent application was conceived no later than August 23, 1994, and was thereafter diligently reduced to practice, both in the United States.
2. I was continuously employed by OPTi Inc. (OPTi) from a time prior to 1994 through a time after 1995.
3. The invention of the above-identified patent application is sometimes referred to as predictive snooping ("pre-snoop") capability, for example of a microprocessor internal cache memory. The patent application describes an embodiment in which, after a PCI-bus controller receives a request from a PCI-bus master to transfer data with an address in secondary memory, the controller performs an initial inquire (snoop) cycle to the microprocessor internal cache. Once the microprocessor signals that the appropriate line of data is either not present in the microprocessor internal cache, or if present is not in a modified state, then the controller allows the burst access to take place

between the memory subsystem and the PCI-bus master. Simultaneously and predictively, the controller also performs an inquire cycle of the microprocessor internal cache for the next cache line. In this manner, if the PCI burst does in fact continue past the cache line boundary, the new inquire cycle will already have taken place (or will already be in progress), thereby allowing the burst to proceed with at most a short delay absent a hit-modified condition.

EXHIBIT 1

4. An embodiment of the invention was designed into an OPTi chipset known as "Viper," specifically the chip designated "82C557 System Controller." Exhibit 1 attached hereto is a true and correct copy of a block diagram, made prior to August 23, 1994, showing how the 82C557 was to be incorporated into a typical computer system. It can be seen that such a system contains a PCI-Bus, to which a PCI device can be connected. A PCI device becomes a master on the PCI-Bus by adhering to a predefined protocol prescribed in the PCI specification.

5. The system of Exhibit 1 also includes a CPU identified as Intel P54C, which has an internal first level cache memory (L1 cache). The system of Exhibit 1 also includes two additional memories identified as Cache and DRAM, which together form part of a memory subsystem. The cache memory shown in Exhibit 1 is a second level cache (L2 cache) that is external to the CPU.

EXHIBIT 2

6. Exhibit 2 appears to be a Purchase Requisition for fabrication of a prototype version of the Viper chipset, including the 82C557 System Controller. This document corroborates my recollection that we sent our design to TSMC, OPTi's fabrication contractor, to fabricate prototype chips from our design in August 1994. From the date on the Purchase Request, it appears to me that we sent the design on or about August 23, 1994.

7. In 1994, the turnaround time for this kind of chip fabrication was on the order of 6-8 weeks. Since during that time, no changes were made to the circuit design of the chips being fabricated, the features included in the prototypes that were delivered to

OPTi in response to this Purchase Request had to have been designed into the chipset no later than August 23, 1994.

EXHIBIT 3

8. Exhibit 3 attached hereto is a true and correct copy of a page from an internal OPTi document entitled "VIPER Applications Training" (hereinafter sometimes referred to as the "Viper Training Manual"). The Viper Training Manual was prepared before the Viper chipset prototypes were delivered to OPTi, and given to, among others, Guarav Shah in the OPTi Applications and Technical Marketing Group for assisting in the debugging of the 82C557 when prototypes arrived. The Viper Training Manual was prepared by members of the Viper design team, including myself and my co-inventor, H.T. Tung.

9. The timing diagram of Exhibit 3 was prepared by Mr. Tung, and corroborates my recollection that predictive snooping was designed into the 82C557 prior to August 23, 1994.

10. The timing diagram of Exhibit 3 is a computer simulation from the logic design of the 82C557, and shows how pre-snoop was designed to operate when a PCI master performs a burst read access to the memory subsystem. In this diagram, all of the data of the burst access is present in the L2 cache (L2 cache hit). Thus the reading of data takes place from the L2 cache rather than from the DRAM memory.

11. On the timing diagram of Exhibit 3, the waveform labeled BVHA represents the address that the 82C557 is driving onto the CPU address bus (HA in Exhibit 1) for the purpose of snooping the CPU's internal L1 cache. The addresses shown in this waveform are expressed in hexadecimal and, reading from left to right and ignoring all but the last three digits, are as follows: 000, 020, 040, 060, 080, and so on. In decimal, these addresses are 0, 32, 64, 92, 128, and so on, respectively. The first of these addresses is the address of a 32-byte line of memory that contains the starting address identified to the 82C557 by a PCI master performing a memory burst read access, and the 82C557 automatically increments this value by 32 bytes in order to generate each subsequent line address. Thus the second memory address that the 82C557 drives onto the CPU address bus is the next sequential line address after the first. Also, if the first

memory address shown in Exhibit 3 is the address of a line L_n of memory, then the second memory address shown in Exhibit 3 is the address of line L_{n+1} .

12. The line address present on the CPU address bus (HA in exhibit 1) is latched for the address input port of the L2 cache by a latch designated "L" in Exhibit 1. The enable input for latch L is illustrated by waveform HACALE in the timing diagram of Exhibit 3. When HACALE is high, the latch is transparent. When HACALE is low, the most recent line address is maintained. Thus from the time that I have labeled A in Exhibit 3 to the time that I have labeled E in Exhibit 3, the data which will be read from the data port of the L2 cache memory will be that of the line containing the starting address identified in the burst read access.

13. Although the actual data port of the L2 cache memory is not shown in the timing diagram of Exhibit 3, the time period during which reading takes place from L2 cache memory can be determined from other waveforms that are shown.

14. The waveforms labeled ECOB and OCOB in Exhibit 3 are the even and odd output enables, respectively, provided to the L2 cache memory. The 82C557 asserts these signals (brings them low) at a time C in Exhibit 3, and they remain asserted until well after time E. Thus from time C to time E, at least, data is being read from the burst read access starting line address in the L2 cache onto the CPU data bus HD (see Exhibit 1).

15. The waveform labeled MDOEX in Exhibit 3 is a control signal that the 82C557 provides to an 82C556, and which, when low, causes the 82C556 to copy data from the HD bus onto the MD bus (see Exhibit 1). As shown in Exhibit 3, MDOEX is low likewise from time C until well after time E. Thus from time C to time E, at least, data is being read from the burst read access starting line address in the L2 cache onto the CPU data bus HD, and then onto the MD bus.

16. The waveform labeled BMDLEX in Exhibit 3 is a control signal that the 82C557 provides to an 82C558 (see Exhibit 1) which controls the transfer of data from the MD bus onto the PCI bus. When BMDLEX is low, data from the MD bus is copied onto the PCI AD bus. When BMDLEX is high, the data most recently copied onto the PCI AD bus is maintained. Since the cache line size in the system of Exhibit 1 is 32 bytes, but the PCI AD bus is only 4 bytes wide (one dword), 8 transfers are required to

transfer an entire line of data onto the PCI AD bus. In the example of Exhibit 3, the burst read access being performed by the PCI master begins with the first dword of a line and continues past the last dword of the line. It can be seen, therefore, that BMDLEX is low 8 times during the period from time C to time E in Exhibit 3 in order to make these 8 transfers. While BMDLEX is high, multiplexers in the 82C556 and 82C558 switch so that the next sequential dword from the HD bus will be selected onto the PCI AD bus the next time BMDLEX is low. Thus from time C to time E, data is being read from the burst read access starting line address in the L2 cache onto the CPU data bus HD, and then onto the MD bus, and is further transferred onto the PCI AD bus.

17. More succinctly, data is being read from the L2 cache memory, at the line address containing the starting address identified in a burst read access, from about time C to about time E in the timing diagram of Exhibit 3.

18. The waveform labeled EADSB in Exhibit 3 shows when inquiry or snoop cycles are being performed on the CPU internal L1 cache. Specifically, EADSB represents the EADS# signal that the 82C557 is providing to the CPU, and indicates a snoop cycle whenever the CPU samples it low. The line address being snooped is the line address that is then present on the CPU HA address lines (represented by the waveform BVHA in Exhibit 3).

19. Thus it can be seen that at a time I have labeled B, an inquiry cycle is being performed on the L1 cache for the line containing the starting address identified in the PCI master's burst read access. It can be seen further that at a time I have labeled D, an inquiry cycle is being performed on the L1 cache for the next sequential cache line after the starting cache line.

20. Moreover, since time D is prior to time E, it can be seen that the 82C557 does not wait for the read data transfers to reach the end of the first cache line before snooping the first level cache for the next sequential cache line. It can also be seen that if the starting address identified by the burst read access is in a cache line L_n , then while data is being read from the L2 cache memory from line L_n according to the burst read access, the 82C557 is simultaneously performing an inquiry cycle of line L_{n+1} in the L1 cache.

TESTING RESULTS

21. After the first prototypes of the 82C557 arrived from OPTi's fabrication contractor, they were given to OPTi's Applications and Technical Marketing Group for testing and debugging. This group, as part of its effort to confirm that all aspects of the chip were working as intended, would attempt to recreate the conditions of the timing diagrams in the Viper Training Manual and monitor the resulting signals to determine whether they conformed to those in the timing diagram. I assisted this group in its efforts to capture the waveforms on a logic analyzer.

22. I specifically recall seeing test results on the logic analyzer, while testing these prototype chips, that satisfied me that predictive snooping was working properly. In particular, I specifically recall seeing waveforms on the logic analyzer in which, similarly to those shown in Exhibit 3, an inquiry cycle was being performed of a line L_{n+1} of the CPU internal L1 cache simultaneously with the reading of data from the memory subsystem at line L_n , where L_n was the address of the line in the memory subsystem containing the starting address of the PCI burst read access. Such waveforms also showed, similarly to those in Exhibit 3, that after the CPU internal L1 cache was snooped for the line containing the starting address of the PCI burst read access, data was transferred from the corresponding line of the memory subsystem according to the burst access. The CPU internal L1 cache was snooped for the next sequential cache line while data from the first line was transferred and without waiting for such transfers to reach the end of the first cache line.

EXHIBIT 4

23. Exhibit 4 is a November 12, 1994 memorandum from Guarav Shah of the OPTi Applications and Technical Marketing Group, to me and others, bearing a "Re:" line of, "82C557 debug status and activity report." I recall receiving a copy of this memorandum.

24. The memorandum indicates that "We have had this silicon for over 2 weeks." The prototype chipset that was the subject of this memorandum therefore had been received by OPTi more than two weeks before the November 12, 1994 date of the memorandum. Given the chip fabrication turnaround time I mentioned above, the date of

this memorandum corroborates my recollection that the version of the chipset that was the subject of this memorandum was the version delivered to OPTi in response to the August 23, 1994 Purchase Request of Exhibit 2. Accordingly, all the features that were reported in the memorandum as being tested, had been designed into the Viper chipset no later than August 23, 1994.

25. As can be seen, this memo includes the following language in the "Behavior" section on the first page:

"Ever since all the tests that we have been doing have the PCI master write wait states set to 2 and there have been no problems. PCI pre-snoop was enabled in both the asynchronous and synchronous modes of operation."

26. This language, together with the date on the memorandum, corroborates my memory that predictive snooping was working properly in the 82C557 prior to November 12, 1994, based on a circuit design that had been completed no later than August 23, 1994.

27. The design and testing of the 82C557 as set forth in this Declaration, as well as the preparation of each of the documents described in this Declaration, all took place in the United States.

28. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that the statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the U.S. Code and that such willful false statements may jeopardize the validity of the application for any patent issued thereon or of any reexamination certificate.

DATED:

Nov 9, 2007



Subir Ghosh

EXHIBIT I TO
SECOND DECLARATION OF SUBIR GHOSH
UNDER 37 C.F.R. §1.131(b)

Viper Block Diagram




EXHIBIT 2 TO
SECOND DECLARATION OF SUBIR GHOSH
UNDER 37 C.F.R. §1.131(b)

SUPPLIER		PURCHASE REQUISITION		OPTI	
SUPPLIER NAME TSMC		THIS IS NOT AN ORDER			
ADDRESS		BUYER NAME		P.O. NUMBER 103464	
CITY		REQUESTOR		DATE 8-23-94	
STATE		PROJECT NAME		DEPARTMENT BIO	
CONTACT PERSON		TELEPHONE NUMBER		ACCOUNT NUMBER	
FAX NUMBER		APPROVAL			

ITEM	SUPPLIER P/N	DESCRIPTION	QTY.	U/M	EST. PRICE	UNIT PRICE	EXTENSION \$	REQ. DELIVERY
1		12 LAYERS FOR PROJECT BRCK6-A			\$	\$	\$	
2		12 LAYERS FOR PROJECT BRCK7-A			\$	\$	\$	
3		super hot lot			\$ on 20# 10	\$ 3480	\$	
4		12 LAYERS FOR PROJECT BRCK8-A			\$	\$	\$	
5		12 LAYERS FOR PROJECT BRCK8-A			\$	\$	\$	
6		8 (Combo) step wafers			\$	\$	\$	
7					\$	\$	\$	
8					\$	\$	\$	
9					\$	\$	\$	
10					\$	\$	\$	
11					\$	\$	\$	

APPROVALS		EST. TOTAL		ACTUAL TOTAL	
PRINT NAME	SIGNATURE	DATE	TELEPHONE NUMBER	FAX NUMBER	CONTACT PERSON
Shirley Yew Tse	<i>[Signature]</i>	8/23/94			
DEPUTY REQUEST MANAGER					
Chia Sun	<i>[Signature]</i>	8/23/94			
APPROVER					
JERRY CHANG	<i>[Signature]</i>				
APPROVED					

FILE: REQPM4	
4-05-94	

EXHIBIT 3 TO
SECOND DECLARATION OF SUBIR GHOSH
UNDER 37 C.F.R. §1.131(b)

EXHIBIT 4 TO
SECOND DECLARATION OF SUBIR GHOSH
UNDER 37 C.F.R. §1.131(b)

RECEIVED

MAR 31 1998

FLIESLER DUBB,
MEYER & LOVEJOY



November 12, 1994

To: Rajesh, HT

Copy: Frank, Sridhar, Subir, Dipankar

From: Gaurav

Re: 82C557 debug status and activity report

Scope:

The purpose of this document is twofold - one is to bring everyone up to speed with the debugging activity on the 82C557, and the other is to serve as a record of how each issue was approached and verified.

Brief:

We have had this silicon for over 2 weeks. We have a CLK trace layout issue in the silicon, due to which we had to set the CAS# precharge to 2 CLKS. That issue prevented us from verifying quite a few issues. For over a week we tried to track down the problem on the silicon and finally realized that it was a violation of a layout rule within the silicon. After that we have had 3 working microsurgery parts on which all the debugging effort has been concentrated.

Discussion:

Each bugfix/enhancement issue that has been tackled has been listed down along with the debugging/verification effort that has been expended on it.

Issue 1:

If the PCI master write wait states have been set to 2 and the PCI bus is running at 33Mhz, then there is a data corruption problem in the L2 cache.

Behavior:

The verification process has been detailed below -

The synchronous PCI mode of operation was chosen. Then the PCI master write waits were set to 2 and the system booted and ran OS/2 for an hour.

The asynchronous PCI mode of operation was chosen. Then the PCI master write waits were set to 2 and the system booted and ran OS/2 for an hour.

Ever since all the tests that we have been doing have the PCI master write wait states set to 2 and there have been no problems. PCI pre-snoop was enabled in both the asynchronous and synchronous modes of operation

We also wrote a small test pattern which was written from the PCI SCSI hard disk to an IDE drive, starting at an odd DoubleWord address. There was no problem with this also.

All the critical signals were hooked up and observed on the logic analyzer. There was no aberration in their behavior.

Issue 2:

If the asynchronous PCI mode of operation is being used, then Index 11h[0] had to be set to a 1. This introduced a one clock delay and prevents L2 cache data corruption.

Behavior:

One does not need to set Index 11h[0] to 1 if an asynchronous PCI mode of operation is being used. The verification process has been detailed below -

Various frequency oscillators were used - 24Mhz, 30Mhz, and 33Mhz - and the system ran OS/2 without any problems.

All the critical signals were hooked up and observed on the logic analyzer. There was no aberration in their behavior.

Issue 3:

6-3-3-3 DRAM timing at 60/66Mhz and 5-3-3-3 timing at 50Mhz

Behavior on 1.3:

This has some problems on the silicon and Subir has an understanding of the failure mechanism.

Issue 4:

Fast NA# generation for single transfer memory write cycles.

Behavior:

This is designed to work in a cacheless (L2) system. It significantly enhances the write performance - close to 35% based on initial test - but it does not work if byte merge has been enabled. The reason for this failure has not yet been investigated.

Issue 5:

Byte Merge fix

Behavior:

There are 3 manifestations of the byte merge problem on this silicon -

- a) If DRAM posted writes has been enabled and if parity checking has been enabled, then enabling byte merge will cause a parity check problem.
- b) If a memory read cycle is piped into a byte merge cycle, then DRAM data corruption takes place due to the wrong assertion of DBCOE1# from the 557.
- c) There are some PCI video cards that will cause the system to hang when byte merge has been enabled.

Issue 6:

Programmable wait states for PCI accesses to DRAM

Untestable features:

The following issues cannot be tested by us without the help of the PCI exerciser.